

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**A SYSTEM AND RELATED METHODS
FOR
REDUCING THE INSTANCES OF SOURCE FILES IN A FILTER
GRAPH**

Inventor(s):
Daniel J. Miller
Eric H. Rudolph

ATTORNEY'S DOCKET NO. MS1-633US

09734529-120600

1 **TECHNICAL FIELD**

2
3 This invention generally relates to processing media content and, more
4 particularly, to an interface and related methods for reducing the instances of
5 source files in a filter graph.
6

7 **BACKGROUND**

8 Recent advances in computing power and related technology have fostered
9 the development of a new generation of powerful software applications. Gaming
10 applications, communications applications, and multimedia applications have
11 particularly benefited from increased processing power and clocking speeds.
12 Indeed, once the province of dedicated, specialty workstations, many personal
13 computing systems now have the capacity to receive, process and render
14 multimedia objects (e.g., audio and video content). While the ability to display
15 (receive, process and render) multimedia content has been around for a while, the
16 ability for a standard computing system to support true multimedia editing
17 applications is relatively new.

18 In an effort to satisfy this need, Microsoft Corporation introduced an
19 innovative development system supporting advanced user-defined multimedia
20 editing functions. An example of this architecture is presented in US Patent No.
21 5,913, 038 issued to Griffiths and commonly owned by the assignee of the present
22 invention, the disclosure of which is expressly incorporated herein by reference.

23 In the '038 patent, Griffiths introduced the an application program interface
24 which, when exposed to higher-level development applications, enable a user to
25 graphically construct a multimedia processing project by piecing together a

collection of “filters” exposed by the interface. The interface described therein is referred to as a filter graph manager. The filter graph manager controls the data structure of the filter graph and the way data moves through the filter graph. The filter graph manager provides a set of component object model (COM) interfaces for communication between a filter graph and its application. Filters of a filter graph architecture are preferably implemented as COM objects, each implementing one or more interfaces, each of which contains a predefined set of functions, called methods. Methods are called by an application program or other component objects in order to communicate with the object exposing the interface. The application program can also call methods or interfaces exposed by the filter graph manager object.

Filter graphs work with data representing a variety of media (or non-media) data types, each type characterized by a data stream that is processed by the filter components comprising the filter graph. A filter positioned closer to the source of the data is referred to as an upstream filter, while those further down the processing chain is referred to as a downstream filter. For each data stream that the filter handles it exposes at least one virtual pin (i.e., distinguished from a physical pin such as one might find on an integrated circuit). A virtual pin can be implemented as a COM object that represents a point of connection for a unidirectional data stream on a filter. Input pins represent inputs and accept data into the filter, while output pins represent outputs and provide data to other filters. Each of the filters include at least one memory buffer, wherein communication of the media stream between filters is accomplished by a series of “copy” operations from one filter to another.

As introduced in Griffiths, a filter graph has three different types of filters: source filters, transform filters, and rendering filters. A source filter is used to load data from some source; a transform filter processes and passes data; and a rendering filter renders data to a hardware device or other locations (e.g., saved to a file, etc.). An example of a filter graph for a simplistic media rendering process is presented with reference to Fig. 1.

Fig. 1 graphically illustrates an example filter graph for rendering media content. As shown, the filter graph 100 is comprised of a plurality of filters 102-114, which read, process (transform) and render media content from a selected source file. As shown, the filter graph includes each of the types of filters described above, interconnected in a linear fashion.

Products utilizing the filter graph have been well received in the market as it has opened the door to multimedia editing using otherwise standard computing systems. It is to be appreciated, however, that the construction and implementation of the filter graphs are computationally intensive and expensive in terms of memory usage. Even the most simple of filter graphs requires and abundance of memory to facilitate the copy operations required to move data between filters. Thus, complex filter graphs can become unwieldy, due in part to the linear nature of conventional development system architecture. Moreover, it is to be appreciated that the filter graphs themselves consume memory resources, thereby compounding the issue introduced above.

Thus, what is required is a filter graph architecture which reduces the computational and memory resources required to support even the most complex of multimedia projects. Just such a solution is disclosed below.

SUMMARY

This invention concerns a system and related interfaces supporting the processing of media content. In accordance with one aspect of the present embodiment, a filter graph for processing media content is presented comprising a video processing subsystem to process video content, and an audio processing subsystem to process audio content. Each of the audio processing subsystem and the video processing subsystem is coupled through a parser to a single instance of a source of audio and video content, wherein the parser selectively provides the audio subsystem and the video subsystem with audio content and video content, respectively. It is to be appreciated that having the audio and video processing subsystems utilize a single parser to access a common source reduces the number of instances of the source file within the filter graph, reducing memory requirements and incrementally improving the perceived performance of the system.

BRIEF DESCRIPTION OF THE DRAWINGS

The same reference numbers are used throughout the figures to reference like components and features.

Fig. 1 is a graphical representation of a conventional filter graph representing a user-defined development project.

Fig. 2 is a block diagram of a computing system incorporating the teachings of the described embodiment.

Fig. 3 is a block diagram of an example software architecture incorporating the teachings of the described embodiment.

Fig. 4 is a graphical illustration of an example software-enabled matrix switch, according to an exemplary embodiment.

Fig. 5 is a graphical representation of a data structure comprising a programming grid to selectively couple one or more of a scalable plurality of input pins to a scalable plurality of output pins of the matrix switch filter, in accordance with one aspect of the described embodiment.

Fig. 6 is a graphical illustration denoting shared buffer memory between filters, according to one aspect of the described embodiment.

Fig. 7 is a flow chart of an example method for generating a filter graph, in accordance with one aspect of the described embodiment.

Fig. 8 is a flow chart of an example method for negotiating buffer requirements between at least two adjacent filters, according to one aspect of the described embodiment.

Fig. 9 graphically illustrates an overview of a process that takes a user-defined editing project and composites a data structure that can be used to program the matrix switch.

Fig. 10 graphically illustrates the project of Fig. 9 in greater detail.

Fig. 11 shows an exemplary matrix switch dynamically generated in support of the project developed in Figs. 9 and 10, according to one described embodiment.

Fig. 12 illustrates a graphic representation of an exemplary data structure that represents the project of Fig. 10, according to one described embodiment.

Figs. 13-18 graphically illustrate various states of a matrix switch programming grid at select points in processing the project of Figs. 9 and 10 through the matrix switch, in accordance with one described embodiment.

Fig. 19 is a flow chart of an example method for processing media content, in accordance with one described embodiment.

Fig. 20 illustrates an example project with a transition and an effect, in accordance with one described embodiment.

Fig. 21 shows an exemplary data structure in the form of a hierarchical tree that represents the project of Fig. 20.

Figs. 22 and 23 graphically illustrate an example matrix switch programming grid associated with the project of Fig. 20 at select points in time, according to one described embodiment.

Fig. 24 shows an example matrix switch dynamically generated and configured as the grid of Figs. 22 and 23 was being processed, in accordance with one described embodiment.

Fig. 25 shows an exemplary project in accordance with one described embodiment.

Fig. 26 graphically illustrates an example audio editing project, according to one described embodiment.

Fig. 27 depicts an example matrix switch programming grid associated with the project of Fig. 26.

Fig. 28 shows an example matrix switch dynamically generated and configured in accordance with the programming grid of Fig. 27 to perform the project of Fig. 26, according to one described embodiment.

Fig. 29 illustrates an exemplary media processing project incorporating another media processing project as a composite, according to yet another described embodiment.

Fig. 30 graphically illustrates an example data structure in the form of a hierarchical tree structure that represents the project of Fig. 29.

Figs 31-36 graphically illustrate various matrix switch programming grid states at select points in generating and configuring the matrix switch to implement the media processing of Fig. 29.

Fig. 38 illustrates an example matrix switch suitable for use in the media processing project of Fig. 29, according to one described embodiment.

Fig. 38a graphically illustrates an example data structure in the form of a hierarchical tree structure that represents a project that is useful in understanding composites in accordance with the described embodiments.

Fig. 39 is a flow diagram that describes steps in a method in accordance with one described embodiment.

Fig. 40 illustrates an example method of generating a filter graph, in accordance with one aspect of the present invention.

Fig. 41 graphically illustrates an example reuse list, according to one aspect of the present invention.

Fig. 42 illustrates an example method for source combining in support of the method introduced in Fig. 40, according to one embodiment of the present invention.

Fig. 43 graphically illustrates a timeline representation of source combining introduced in Fig. 42.

Fig. 44 illustrates a block diagram of an example render engine utilizing a segment filter, in accordance with one aspect of the present invention.

Fig. 45 illustrates a flow chart of an example method of generating a filter graph to reuse source filters, in accordance with one aspect of the present invention.

Fig. 46 illustrates a flow chart of an example method of executing a development project utilizing a segment filter, in accordance with one aspect of the present invention.

Fig. 47 illustrates a block diagram of an example media parser filter, according to one embodiment of the present invention.

Fig. 48 illustrates a flow chart of an example method for building a filter graph, according to one aspect of the present invention.

Fig. 49 graphically illustrates an example filter graph sharing a file parser between different media processing sub-systems, according to one embodiment of the present invention.

DETAILED DESCRIPTION

Related Applications

This application is related to the following commonly-filed U.S. Patent Applications, all of which are commonly assigned to Microsoft Corp., the disclosures of which are incorporated by reference herein:

- Application Serial No. _____, entitled "An Interface and Related Methods for Reducing Source Accesses in a Development System", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-643US;
- Application Serial No. _____, entitled "A System and Related Interfaces Supporting the Processing of Media Content", naming

Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-629US;

- Application Serial No. _____, entitled "A System and Related Methods for Reducing Source Filter Invocation in a Development Project", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-631US;
- Application Serial No. _____, entitled "A System and Related Methods for Reducing Memory Requirements of a Media Processing System", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-632US;
- Application Serial No. _____, entitled "An Interface and Related Methods for Dynamically Generating a Filter Graph in a Development System", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-634US;
- Application Serial No. _____, entitled "A System and Related Methods for Processing Audio Content in a Filter Graph", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-639US;
- Application Serial No. _____, entitled "A System and Methods for Generating an Managing Filter Strings in a Filter Graph", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-642US;
- Application Serial No. _____, entitled "Methods and Systems for Processing Media Content", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-640US;
- Application Serial No. _____, entitled "Systems for Managing Multiple Inputs and Methods and Systems for Processing Media Content ", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-635US;
- Application Serial No. _____, entitled "Methods and Systems for Implementing Dynamic Properties on Objects that Support Only Static Properties", naming Daniel J. Miller and David Maymudes as inventors, and bearing attorney docket number MS1-638US;
- Application Serial No. _____, entitled "Methods and Systems for Efficiently Processing Compressed and Uncompressed Media Content", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-630US;
- Application Serial No. _____, entitled "Methods and Systems for Effecting Video Transitions Represented By Bitmaps", naming Daniel J. Miller and David Maymudes as inventors, and bearing attorney docket number MS1-637US;

- Application Serial No. _____, entitled “Methods and Systems for Mixing Digital Audio Signals”, naming Eric H. Rudolph as inventor, and bearing attorney docket number MS1-636US; and
- Application Serial No. _____, entitled “Methods and Systems for Processing Multi-media Editing Projects”, naming Eric H. Rudolph as inventor, and bearing attorney docket number MS1-641US.

Various described embodiments concern an application program interface associated with a development system. According to one example implementation, the interface is exposed to a media processing application to enable a user to dynamically generate complex media processing tasks, e.g., editing projects. In the discussion herein, aspects of the invention are developed within the general context of computer-executable instructions, such as program modules, being executed by one or more conventional computers. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including hand-held devices, personal digital assistants, multiprocessor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, and the like. In a distributed computer environment, program modules may be located in both local and remote memory storage devices. It is noted, however, that modification to the architecture and methods described herein may well be made without deviating from spirit and scope of the present invention. Moreover, although developed within the context of a media processing system paradigm, those skilled in the art will appreciate, from the discussion to follow, that the application program interface may well be applied to other development

1 system implementations. Thus, the media processing system described below is
2 but one illustrative implementation of a broader inventive concept.

3 4 **Example System Architecture**

5 **Fig. 2** illustrates an example of a suitable computing environment 200 on
6 which the system and related methods for processing media content may be
7 implemented.

8 It is to be appreciated that computing environment 200 is only one example
9 of a suitable computing environment and is not intended to suggest any limitation
10 as to the scope of use or functionality of the media processing system. Neither
11 should the computing environment 200 be interpreted as having any dependency
12 or requirement relating to any one or combination of components illustrated in the
13 exemplary computing environment 200.

14 The media processing system is operational with numerous other general
15 purpose or special purpose computing system environments or configurations.
16 Examples of well known computing systems, environments, and/or configurations
17 that may be suitable for use with the media processing system include, but are not
18 limited to, personal computers, server computers, thin clients, thick clients, hand-
19 held or laptop devices, multiprocessor systems, microprocessor-based systems, set
20 top boxes, programmable consumer electronics, network PCs, minicomputers,
21 mainframe computers, distributed computing environments that include any of the
22 above systems or devices, and the like.

23 In certain implementations, the system and related methods for processing
24 media content may well be described in the general context of computer-
25 executable instructions, such as program modules, being executed by a computer.

1 Generally, program modules include routines, programs, objects, components,
2 data structures, etc. that perform particular tasks or implement particular abstract
3 data types. The media processing system may also be practiced in distributed
4 computing environments where tasks are performed by remote processing devices
5 that are linked through a communications network. In a distributed computing
6 environment, program modules may be located in both local and remote computer
7 storage media including memory storage devices.

8 In accordance with the illustrated example embodiment of Fig. 2 computing
9 system 200 is shown comprising one or more processors or processing units 202, a
10 system memory 204, and a bus 206 that couples various system components
11 including the system memory 204 to the processor 202.

12 Bus 206 is intended to represent one or more of any of several types of bus
13 structures, including a memory bus or memory controller, a peripheral bus, an
14 accelerated graphics port, and a processor or local bus using any of a variety of
15 bus architectures. By way of example, and not limitation, such architectures
16 include Industry Standard Architecture (ISA) bus, Micro Channel Architecture
17 (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association
18 (VESA) local bus, and Peripheral Component Interconnects (PCI) buss also
19 known as Mezzanine bus.

20 Computer 200 typically includes a variety of computer readable media.
21 Such media may be any available media that is locally and/or remotely accessible
22 by computer 200, and it includes both volatile and non-volatile media, removable
23 and non-removable media.

24 In Fig. 2, the system memory 204 includes computer readable media in the
25 form of volatile, such as random access memory (RAM) 210, and/or non-volatile

memory, such as read only memory (ROM) 208. A basic input/output system (BIOS) 212, containing the basic routines that help to transfer information between elements within computer 200, such as during start-up, is stored in ROM 208. RAM 210 typically contains data and/or program modules that are immediately accessible to and/or presently be operated on by processing unit(s) 202.

Computer 200 may further include other removable/non-removable, volatile/non-volatile computer storage media. By way of example only, Fig. 2 illustrates a hard disk drive 228 for reading from and writing to a non-removable, non-volatile magnetic media (not shown and typically called a "hard drive"), a magnetic disk drive 230 for reading from and writing to a removable, non-volatile magnetic disk 232 (e.g., a "floppy disk"), and an optical disk drive 234 for reading from or writing to a removable, non-volatile optical disk 236 such as a CD-ROM, DVD-ROM or other optical media. The hard disk drive 228, magnetic disk drive 230, and optical disk drive 234 are each connected to bus 206 by one or more interfaces 226.

The drives and their associated computer-readable media provide nonvolatile storage of computer readable instructions, data structures, program modules, and other data for computer 200. Although the exemplary environment described herein employs a hard disk 228, a removable magnetic disk 232 and a removable optical disk 236, it should be appreciated by those skilled in the art that other types of computer readable media which can store data that is accessible by a computer, such as magnetic cassettes, flash memory cards, digital video disks, random access memories (RAMs), read only memories (ROM), and the like, may also be used in the exemplary operating environment.

1 A number of program modules may be stored on the hard disk 228,
2 magnetic disk 232, optical disk 236, ROM 208, or RAM 210, including, by way of
3 example, and not limitation, an operating system 214, one or more application
4 programs 216 (e.g., multimedia application program 224), other program modules
5 218, and program data 220. In accordance with the illustrated example
6 embodiment of Fig. 2, operating system 214 includes an application program
7 interface embodied as a render engine 222. As will be developed more fully
8 below, render engine 222 is exposed to higher-level applications (e.g., 216) to
9 automatically assemble filter graphs in support of user-defined development
10 projects, e.g., media processing projects. Unlike conventional media processing
11 systems, however, render engine 222 utilizes a scalable, dynamically
12 reconfigurable matrix switch to reduce filter graph complexity, thereby reducing
13 the computational and memory resources required to complete a development
14 project. Various aspects of the innovative media processing system represented by
15 a computer 200 implementing the innovative render engine 222 will be developed
16 further, below.

17 Continuing with Fig. 2, a user may enter commands and information into
18 computer 200 through input devices such as keyboard 238 and pointing device 240
19 (such as a "mouse"). Other input devices may include a audio/video input
20 device(s) 253, a microphone, joystick, game pad, satellite dish, serial port, scanner,
21 or the like (not shown). These and other input devices are connected to the
22 processing unit(s) 202 through input interface(s) 242 that is coupled to bus 206,
23 but may be connected by other interface and bus structures, such as a parallel port,
24 game port, or a universal serial bus (USB).
25

1 A monitor 256 or other type of display device is also connected to bus 206
2 via an interface, such as a video adapter 244. In addition to the monitor, personal
3 computers typically include other peripheral output devices (not shown), such as
4 speakers and printers, which may be connected through output peripheral interface
5 246.

6 Computer 200 may operate in a networked environment using logical
7 connections to one or more remote computers, such as a remote computer 250.
8 Remote computer 250 may include many or all of the elements and features
9 described herein relative to computer 200 including, for example, render engine
10 222 and one or more development applications 216 utilizing the resources of
11 render engine 222.

12 As shown in Fig. 2, computing system 200 is communicatively coupled to
13 remote devices (e.g., remote computer 250) through a local area network (LAN)
14 251 and a general wide area network (WAN) 252. Such networking environments
15 are commonplace in offices, enterprise-wide computer networks, intranets, and the
16 Internet.

17 When used in a LAN networking environment, the computer 200 is
18 connected to LAN 251 through a suitable network interface or adapter 248. When
19 used in a WAN networking environment, the computer 200 typically includes a
20 modem 254 or other means for establishing communications over the WAN 252.
21 The modem 254, which may be internal or external, may be connected to the
22 system bus 206 via the user input interface 242, or other appropriate mechanism.

23 In a networked environment, program modules depicted relative to the
24 personal computer 200, or portions thereof, may be stored in a remote memory
25 storage device. By way of example, and not limitation, Fig. 2 illustrates remote

1 application programs 216 as residing on a memory device of remote computer
2 250. It will be appreciated that the network connections shown and described are
3 exemplary and other means of establishing a communications link between the
4 computers may be used.

5 Turning next to **Fig. 3**, a block diagram of an example development system
6 architecture is presented, in accordance with one embodiment of the present
7 invention. In accordance with the illustrated example embodiment of Fig. 3,
8 development system 300 is shown comprising one or more application program(s)
9 216 coupled to render engine 222 via an appropriate communications interface
10 302. As used herein, application program(s) 216 are intended to represent any of a
11 wide variety of applications which may benefit from use of render engine 222
12 such as, for example a media processing application 224.

13 The communications interface 302 is intended to represent any of a number
14 of alternate interfaces used by operating systems to expose application program
15 interface(s) to applications. According to one example implementation, interface
16 302 is a component object model (COM) interface, as used by operating systems
17 offered by Microsoft Corporation. As introduced above, COM interface 302
18 provides a means by which the features of the render engine 222, to be described
19 more fully below, are exposed to an application program 216.

20 In accordance with the illustrated example implementation of Fig. 3, render
21 engine 222 is presented comprising source filter(s) 304A-N, transform filter(s)
22 306A-N and render filter 310, coupled together utilizing virtual pins to facilitate a
23 user-defined media processing project. According to one implementation, the
24 filters of system 300 are similar to the filters exposed in conventional media
25 processing systems. According to one implementation, however, filters are not

coupled via such interface pins. Rather, alternate implementations are envisioned wherein individual filters (implemented as objects) make calls to other objects, under the control of the render engine 222, for the desired input. Unlike conventional systems, however, render engine 222 exposes a scalable, dynamically reconfigurable matrix switch filter 308, automatically generated and dynamically configured by render engine 222 to reduce the computational and memory resource requirements often associated with development projects. As introduced above, the pins (input and/or output) are application interface(s) designed to communicatively couple other objects (e.g., filters).

In accordance with the example implementation of a media processing system, an application communicates with an instance of render engine 222 when the application 216 wants to process streaming media content. Render engine 222 selectively invokes and controls an instance of filter graph manager (not shown) to automatically create a filter graph by invoking the appropriate filters (e.g., source, transform and rendering). As introduced above, the communication of media content between filters is achieved by either (1) coupling virtual output pins of one filter to the virtual input pins of requesting filter; or (2) by scheduling object calls between appropriate filters to communicate the requested information. As shown, source filter 304 receives streaming data from the invoking application or an external source (not shown). It is to be appreciated that the streaming data can be obtained from a file on a disk, a network, a satellite feed, an Internet server, a video cassette recorder, or other source of media content. As introduced above, transform filter(s) 306 take the media content and processes it in some manner, before passing it along to render filter 310. As used herein, transform filter(s) 306 are intended to represent a wide variety of processing methods or applications that

1 can be performed on media content. In this regard, transform filter(s) 306 may
2 well include a splitter, a decoder, a sizing filter, a transition filter, an effects filter,
3 and the like. The function of each of these filters is described more fully in the
4 Griffiths application, introduced above, and generally incorporated herein by
5 reference. The transition filter, as used herein, is utilized by render engine 222 to
6 transition the rendered output from a first source to a second source. The effect
7 filter is selectively invoked to introduce a particular effect (e.g., fade, wipe, audio
8 distortion, etc.) to a media stream.

9 In accordance with one aspect of the embodiment, to be described more
10 fully below, matrix switch filter 308 selectively passes media content from one or
11 more of a scalable plurality of input(s) to a scalable plurality of output(s).
12 Moreover, matrix switch 308 also supports implementation of a cascaded
13 architecture utilizing feedback paths, i.e., wherein transform filters 306B, 306C,
14 etc. coupled to the output of matrix switch 308 are dynamically coupled to one or
15 more of the scalable plurality of matrix switch input(s). An example of this
16 cascaded filter graph architecture is introduced in Fig. 3, and further explained in
17 example implementations, below.

18 Typically, media processed through source, transform and matrix switch
19 filters are ultimately passed to render filter 310, which provides the necessary
20 interface to a hardware device, or other location that accepts the renderer output
21 format, such as a memory or disk file, or a rendering device.

22 **Fig. 4** is a graphical illustration of an example software-enabled matrix
23 switch 308, according to one example embodiment of the present invention. As
24 shown, the matrix switch 308 is comprised of a scalable plurality of input(s) 402
25 and a scalable plurality of output(s) 404, wherein any one or more of the input(s)

402 may be iteratively coupled to any one or more of the output(s) 404, based on the content of the matrix switch programming grid 406, automatically generated by render engine 222. According to an alternate implementation introduced above, switch matrix 308 is programmed by render engine 222 to dynamically generate object calls to communicate media content between filters. In addition, according to one implementation, matrix switch 308 includes a plurality of input/output (I/O) buffers 408, as well as means for maintaining source, or media time 410 and/or timeline, or project time 412. It is to be appreciated, however, that in alternate implementations matrix switch 308 does not maintain both source and project times, relying on an upstream filter to convert between these times. As will be developed more fully below, matrix switch 308 dynamically couples one or more of the scalable plurality of inputs 402 to one or more of the scalable plurality of outputs 404 based, at least in part, on the media time 410 and/or the project time 412 and the content of matrix switch programming grid 406. In this regard, matrix switch 308 may be characterized as time-aware, supporting such advanced editing features as searching/seeking to a particular point (e.g., media time) in the media content, facilitating an innovative buffering process utilizing I/O buffers 408 to facilitate look-ahead processing of media content, and the like. Thus, it will be appreciated given the discussion to follow that introduction of the matrix switch 308 provides a user with an editing flexibility that was heretofore unavailable in a personal computer-based media processing system.

As introduced above, the inputs 402 and outputs 404 of matrix switch 308 are interfaces which facilitate the time-sensitive routing of data (e.g., media content) in accordance with a user-defined development project. Matrix switch 308 has a scalable plurality of inputs 402 and outputs 404, meaning that the

number of inputs 402 and outputs 404 are individually generated to satisfy a given editing project. Insofar as each of the inputs/outputs (I/O) has an associated transfer buffer (preferably shared with an adjacent filter) to communicate media content, the scalability of the input/output serves to reduce the overall buffer memory consumed by an editing project. According to one implementation, output 1 is generally reserved as a primary output, e.g., coupled to a rendering filter (not shown).

According to one implementation, for each input 402 and output 404, matrix switch 308 attempts to be the allocator, or manager of the buffer associated with the I/O(s) shared with adjacent filters. One reason is to ensure that all of the buffers are of the same size and share common attributes so that a buffer associated with any input 402 may be shared with any output 404, thereby reducing the need to copy memory contents between individual buffers associated with such inputs/outputs. If matrix switch 308 cannot be an allocator for a given output (404), communication from an input (402) to that output is performed using a conventional memory copy operation between the individual buffers associated with the select input/output.

As introduced above, the matrix switch programming grid 406 is dynamically generated by render engine 222 based, at least in part, on the user-defined development project. As will be developed below, render engine 222 invokes an instance of filter graph manager to assemble a tree structure of an editing project, noting dependencies between source, filters and time to dynamically generate the programming grid 406. A data structure comprising an example programming grid 406 is introduced with reference to Fig. 5, below.

0974339-120600

Turning briefly to **Fig. 5**, a graphical representation of a data structure comprising an example programming grid 406 is presented, in accordance with one embodiment of the present invention. In accordance with the illustrated example embodiment of Fig. 5, programming grid 406 is depicted as a two-dimensional data structure comprising a column along the y-axis 502 of the grid denoting input pins associated with a content chain (e.g., series of filters to process media content) of the development project. The top row along the x-axis 504 of the data structure denotes project time. With these grid “borders”, the body 506 of the grid 406 is populated with output pin assignments, denoting which input pin is coupled to which output pin during execution of the development project. In this way, render engine 222 dynamically generates and facilitates matrix switch 308. Those skilled in the art will appreciate, however, that data structures of greater or lesser complexity may well be used in support of the programming grid 406 without deviating from the spirit and scope of the present invention.

Returning to Fig. 4, matrix switch 308 is also depicted with a plurality of input/output buffers 408, shared among all of the input(s)/output(s) (402, 404) to facilitate advanced processing features. That is, while not required to implement the core features of matrix switch 308, I/O buffers 408 facilitate a number of innovative performance enhancing features to improve the performance (or at least the user’s perception of performance) of the processing system, thereby providing an improved user experience. According to one implementation, I/O buffers 408 are separate from the buffers assigned to each individual input and output pin in support of communication through the switch. According to one implementation, I/O buffers 408 are primarily used to foster look-ahead processing of the project. Assume, for example, that a large portion of the media processing project required

only 50% of the available processing power, while some smaller portion required 150% of the available processing power. Implementation of the shared I/O buffers 408 enable filter graph manager to execute tasks ahead of schedule and buffer this content in the shared I/O buffers 408 until required. Thus, when execution of the filter graph reaches a point where more than 100% of the available processing power is required, the processing system can continue to supply content from the I/O buffers 408, while the system completes execution of the CPU-intensive tasks. If enough shared buffer space is provided, the user should never know that some tasks were not performed in real-time. According to one implementation, shared buffers 408 are dynamically split into two groups by render engine 222, a first group supports the input(s) 402, while a second (often smaller) group is used in support of a primary output (e.g., output pin 1) to facilitate a second, independent output processing thread. The use of an independent output buffers the render engine from processing delays that might occur in upstream and/or downstream filters, as discussed above. It will be appreciated by those skilled in the art that such that matrix switch 308 and the foregoing described architecture beneficially suited to support media streaming applications.

As introduced above, the filter graph is time-aware in the sense that media (source) time and project execution time are maintained. According to one implementation, matrix switch 308 maintains at least the project clock, while an upstream filter maintains the source time, converting between source and project time for all downstream filters (i.e., including the matrix switch 308). According to one implementation, the frame rate converter filter of a filter graph is responsible for converting source time to project time, and vice versa, i.e.,

supporting random seeks, etc. Alternatively, matrix switch 308 utilizes an integrated set of clock(s) to independently maintain project and media times.

Having introduced the architectural and operational elements of matrix switch filter 308, **Fig. 6** graphically illustrates an example filter graph implementation incorporating the innovative matrix switch 308. In accordance with the illustrated example embodiment, filter graph 600 is generated by render engine 222 in response to a user defined development project. Unlike the lengthy linear filter graphs typical of convention development systems however, filter graph 600 is shown incorporating a matrix switch filter 308 to recursively route the pre-processed content (e.g., through filters 602, 606, 610, 614 and 618, described more fully below) through a user-defined number of transform filters including, for example, transition filter(s) 620 and effects filter(s) 622. Moreover, as will be developed more fully below, the scalable nature of matrix switch filter 308 facilitates such iterative processing for any number of content threads, tracks or compositions.

According to one implementation, a matrix switch filter 308 can only process one type of media content, of the same size and at the same frame-rate (video) or modulation type/schema (audio). Thus, Fig. 6 is depicted comprising pre-processing filters with a parser filter 606 to separate, independent content type(s) (e.g., audio content and video content), wherein one of the media types would be processed along a different path including a separate instance of matrix switch 308. Thus, in accordance with the illustrated example embodiment of a media processing system, processing multimedia content including audio and video would utilize two (2) matrix switch filters 308, one dedicated to audio processing (not shown) and one dedicated to video processing. That is not to say,

1 however, that multiple switch filters 308 could not be used (e.g., two each for
2 audio and video) for each content type in alternate implementations. Similarly, it
3 is anticipated that in alternate implementations a matrix switch 308 that accepts
4 multiple media types could well be used without deviating from the spirit and
5 scope of the present invention.

6 In addition filter graph 600 includes a decoder filter 610 to decode the
7 media content. Resize filter 614 is employed when matrix switch 308 is to receive
8 content from multiple sources, ensuring that the size of the received content is the
9 same, regardless of the source. According to one implementation, resize filter 614
10 is selectively employed in video processing paths to adjust the media size of
11 content from one or more sources to a user-defined level. Alternatively, resizer
12 filter 614 adjusts the media size to the largest size provided by any one or more
13 media sources. That is, if, for example, render engine 222 identifies the largest
14 required media size (e.g., 1270x1040 video pixels per frame) and, for any content
15 source not providing content at this size, the content is modified (e.g., stretched,
16 packed, etc.) to fill this size requirement. The frame rate converter (FRC) and
17 pack filter 618, introduced above, ensures that video content from the multiple
18 sources is arriving at the same frame rate, e.g., ten (10) frames per second. As
19 introduced above, the FRC also maintains the distinction between source time and
20 project time.

21 In accordance with one aspect of the present invention, filter graph 600 is
22 depicted utilizing a single, negotiated buffer 604, 608, 612, 616, etc. between
23 adjacent filters. In this regard, render engine 222 reduces the buffer memory
24 requirements in support of a development project.
25

From the point of pre-processing (filters 602, 606, 610, 614, 618), rather than continue a linear filter graph incorporating all of the transition 620 and effect 622 filter(s), render engine 222 utilizes a cascade architecture, recursively passing media content through the matrix switch 308 to apply to the transform filter(s) (e.g., 620, 622, etc.) to complete the execution of the development project. It will be appreciated by those skilled in the art that the ability to recursively pass media content through one or more effect and/or transition filters provided by the matrix switch filter 308 greatly reduces the perceived complexity of otherwise large filter graphs, while reducing memory and computational overhead.

Turning to **Fig. 7**, a flow chart of an example method for generating a filter graph is presented, in accordance with one aspect of the present invention. The method 700 begins with block 702 wherein render engine 222 receives an indication to generate a filter graph representing a user-defined development project (e.g., a media editing project). According to one example implementation, the indication is received from an application 224 via COM interface(s) 302.

In block 704, render engine 222 facilitates generation of the editing project, identifying the number and type of media sources selected by the user. In block 706, based at least in part on the number and/or type of media sources, filter graph manger 222 exposes source, transform and rendering filter(s) to effect a user defined media processing project, while beginning to establish a programming grid 406 for the matrix switch filter 308.

In block 708, reflecting user editing instructions, render engine 222 completes the programming grid 406 for matrix switch 308, identifying which inputs 402 are to be coupled to which outputs 404 at particular project times.

Based, at least in part, on the programming grid 406 render engine 222 generates a matrix switch filter 308 with an appropriate number of input 402 and output 404 pins to effect the project, and assembles the filter graph, block 710.

In block 712, to reduce the buffer memory requirements for the processing project, the render engine 222 instructs the filters populating the filter graph to (re)negotiate buffer memory requirements between filters. That is, adjacent filters attempt to negotiate a size and attribute standard so that a single buffer can be utilized to couple each an output pin of one filter to an input pin of a downstream filter. An example implementation of the buffer negotiation process of block 712 is presented in greater detail with reference to Fig. 8.

Turning briefly to Fig. 8, an example method of negotiating buffer requirements between adjacent filters is presented, in accordance with one example implementation of the present invention. Once the final connection is established to matrix switch 308, matrix switch 308 identifies the maximum buffer requirements for any filter coupled to any of its pins (input 402 and/or output 404), block 802. According to one implementation, the maximum buffer requirements are defined as the lowest common multiple of buffer alignment requirements, and the maximum of all the pre-fix requirements of the filter buffers.

In block 804, matrix switch 308 selectively removes one or more existing filter connections to adjacent filters. Matrix switch 308 then reconnects all of its pins to adjacent filters using a common buffer size between each of the pins, block 806. In block 808, matrix switch 308 negotiates to be the allocator for all of its pins (402, 404). If the matrix switch 308 cannot, for whatever reason, be the allocator for any of its input pins 402 minimal loss to performance is encountered, as the buffer associated with the input pin will still be compatible with any

1 downstream filter (i.e., coupled to an output pin) and, thus, the buffer can still be
2 passed to the downstream filter without requiring a memory copy operation. If,
3 however, matrix switch 308 cannot be an allocator for one of its output pins 404,
4 media content must then be transferred to at least the downstream filter associated
5 with that output pin using a memory copy operation, block 810.

6 In block 812, once the matrix switch 308 has re-established its connection
7 to adjacent filters, render engine 222 restores the connection in remaining filters
8 using negotiated buffer requirements emanating from the matrix switch filter 308
9 buffer negotiations. Once the connections throughout the filter graph have been
10 reconnected, the process continues with block 714 of Fig. 7.

11 In block 714 (Fig. 7), have re-established the connections between filters,
12 render engine 222 is ready to implement a user's instruction to execute the media
13 processing project.

14 15 **Example Operation and Implementation(s)**

16 The matrix switch described above is quite useful in that it allows multiple
17 inputs to be directed to multiple outputs at any one time. These input can compete
18 for a matrix switch output. The embodiments described below permit these
19 competing inputs to be organized so that the inputs smoothly flow through the
20 matrix switch to provide a desired output. And, while the inventive programming
21 techniques are described in connection with the matrix switch as such is employed
22 in the context of multi-media editing projects, it should be clearly understood that
23 application of the inventive programming techniques and structures should not be
24 so limited only to application in the field of multi-media editing projects or, for
25 that matter, multi-media applications or data streams. Accordingly, the principles

1 about to be discussed can be applied to other fields of endeavor in which multiple
2 inputs can be characterized as competing for a particular output during a common
3 time period.

4 In the multi-media example below, the primary output of the matrix switch
5 is a data stream that defines an editing project that has been created by a user.
6 Recall that this editing project can include multiple different sources that are
7 combined in any number of different ways, and the sources that make up a project
8 can comprise audio sources, video sources, or both. The organization of the inputs
9 and outputs of the matrix switch are made manageable, in the examples described
10 below, by a data structure that permits the matrix switch to be programmed.

11 Fig. 9 shows an overview of a process that takes a user-defined editing
12 project and renders from it a data structure that can be used to program the matrix
13 switch.

14 Specifically, a user-defined editing project is shown generally at 900.
15 Typically, when a user creates an editing project, they can select from a number of
16 different multimedia clips that they can then assemble into a unique presentation.
17 Each individual clip represents a *source* of digital data or a source stream (e.g.,
18 multimedia content). Projects can include one or more sources 902. In defining
19 their project, a user can operate on sources in different ways. For example, video
20 sources can have *transitions* 904 and *effects* 906 applied on them. A transition
21 object is a way to change between two or more sources. As discussed above, a
22 transition essentially receives as input, two or more streams, operates on them in
23 some way, and produces a single output stream. An exemplary transition can
24 comprise, for example, fading from one source to another. An effect object can
25 operate on a single source or on a composite of sources. An effect essentially

0974329-136600

1 receives a single input stream, operates on it in some way, and produces a single
2 output stream. An exemplary effect can comprise a black-and-white effect in
3 which a video stream that is configured for presentation in color format is
4 rendered into a video stream that is configured for presentation in black and white
5 format. Unlike conventional effect filters, effect object 906 may well perform
6 multiple effect tasks. That is, in accordance with one implementation, an effect
7 object (e.g., 906) may actually perform multiple tasks on the received input
8 stream, wherein said tasks would require multiple effect filters in a conventional
9 filter graph system.

10 An exemplary user interface 908 is shown and represents what a user might
11 see when they produce a multimedia project with software executing on a
12 computer. In this example, the user has selected three sources A, B, and C, and
13 has assembled the sources into a project timeline. The project timeline defines
14 when the individual sources are to be rendered, as well as when any transitions
15 and/or effects are to occur.

16 In the discussion that follows, the notion of a *track* is introduced. A track
17 can contain one or more sources or source clips. If a track contains more than one
18 source clip, the source clips cannot overlap. If source clips are to overlap (e.g.
19 fading from one source to another, or having one source obscure another), then
20 multiple tracks are used. A track can thus logically represent a layer on which
21 sequential video is produced. User interface 908 illustrates a project that utilizes
22 three tracks, each of which contains a different source. In this particular project
23 source A will show for a period of time. At a defined time in the presentation,
24 source A is obscured by source B. At some later time, source B transitions to
25 source C.

1 In accordance with the described embodiment, the user-defined editing
2 project 900 is translated into a data structure 910 that represents the project. In the
3 illustrated and described example, this data structure 910 comprises a tree
4 structure. It is to be understood, however, that other data structures could be used.
5 The use of tree structures to represent editing projects is well-known and is not
6 described here in any additional detail. Once the data structure 910 is defined, it is
7 processed to provide a data structure 912 that is utilized to program the matrix
8 switch. In the illustrated and described embodiment, data structure 912 comprises
9 a grid from which the matrix switch can be programmed. It is to be understood
10 and appreciated that other data structures and techniques could, however, be used
11 to program the matrix switch without departing from the spirit and scope of the
12 claimed subject matter.

13 The processing that takes place to define data structures 910 and 912 can
14 take place using any suitable hardware, software, firmware, or combination
15 thereof. In the examples set forth below, the processing takes place utilizing
16 software in the form of a video editing software package that is executable on a
17 general purpose computer.

18 19 Example Project

20 For purposes of explanation, consider Fig. 10 which shows project 908
21 from Fig. 9 in a little additional detail. Here, a time line containing numbers 0-16
22 is provided adjacent the project to indicate when particular sources are to be seen
23 and when transitions and effects (when present) are to occur. In the examples in
24 this document, the following convention exists with respect to projects, such as
25 project 908. A priority exists for video portions of the project such that as one

proceeds from top to bottom, the priority increases. Thus, in the Fig. 10 example, source A has the lowest priority followed by source B and source C. Thus, if there is an overlap between higher and lower priority sources, the higher priority source will prevail. For example, source B will obscure source A from between $t = 4-8$.

In this example, the following can be ascertained from the project 908 and time line: from time $t=0-4$ source A should be routed to the matrix switch's primary output; from $t=4-12$ source B should be routed to the matrix switch's primary output; from $t=12-14$ there should be a transition between source B and source C which should be routed to the matrix switch's primary output; and from $t=14-16$ source C should be routed to the matrix switch's primary output. Thus, relative to the matrix switch, each of the sources and the transition can be characterized by where it is to be routed at any given time. Consider, for example, the table just below:

Object	Routing for a given time
C	$t = 0-12$ (nowhere); $t = 12-14$ (transition); $t = 14-16$ (primary output)
B	$t = 0-4$ (nowhere); $t = 4-12$ (primary output); $t = 12-14$ (transition); $t = 14-16$ (nowhere)
A	$t = 0-4$ (primary output); $t = 4-16$ (nowhere)
Transition	$t = 0-12$ (nowhere); $t = 12-14$ (primary output); $t = 14-16$ (nowhere)

Fig. 11 shows an exemplary matrix switch 1100 that can be utilized in the presentation of the user's project. Matrix switch 1100 comprises multiple inputs and multiple outputs. Recall that a characteristic of the matrix switch 1100 is that any of the inputs can be routed to any of the outputs at any given time. A

1 transition element 1102 is provided and represents the transition that is to occur
2 between sources B and C. Notice that the matrix switch includes four inputs
3 numbered 0-3 and three outputs numbered 0-2. Inputs 0-2 correspond respectively
4 to sources A-C, while input 3 corresponds to the output of the transition element
5 1102. Output 0 corresponds to the switch's primary output, while outputs 1 and 2
6 are routed to the transition element 1102.

7 The information that is contained in the table above is the information that
8 is utilized to program the matrix switch. The discussion presented below describes
9 but one implementation in which the information contained in the above table can
10 be derived from the user's project time line.

11 Recall that as a user edits or creates a project, software that comprises a part
12 of their editing software builds a data structure that represents the project. In the
13 Fig. 9 overview, this was data structure 910. In addition to building the data
14 structure that represents the editing project, the software also builds and configures
15 a matrix switch that is to be used to define the output stream that embodies the
16 project. Building and configuring the matrix switch can include building the
17 appropriate graphs (e.g., a collection of software objects, or filters) that are
18 associated with each of the sources and associating those graphs with the correct
19 inputs of the matrix switch. In addition, building and configuring the matrix
20 switch can also include obtaining and incorporating additional appropriate filters
21 with the matrix switch, e.g. filters for transitions, effects, and mixing (for audio
22 streams). This will become more apparent below.

23 Fig. 12 shows a graphic representation of an exemplary data structure 1200
24 that represents the project of Fig. 10. Here, the data structure comprises a
25 traditional hierarchical tree structure. Any suitable data structure can, however, be

1 particular node's object desires to be routed to the output, e.g. node 1204a would
2 include an identifier for the node associating it with source A, as well as a time
3 value that indicates that source A desires to be routed to the output from time $t = 0$ -
4 8. This meta information is utilized to build the data structure that is, in turn,
5 utilized to program the matrix switch.

6 In the example about to be described below, a specific data structure in the
7 form of a grid is utilized. In addition, certain specifics are described with respect
8 to how the grid is processed so that the matrix switch can be programmed. It is to
9 be understood that the specific described approach is for exemplary purposes only
10 and is not intended to limit application of the claims. Rather, the specific approach
11 constitutes but one way of implementing broader conceptual notions embodied by
12 the inventive subject matter.

13 Figs. 13-18 represent a process through which the inventive grid is built. In
14 the grid about to be described, the x axis represents time, and the y axis represents
15 layers in terms of priority that go from lowest (at the top of the grid) to highest (at
16 the bottom of the grid). Every row in the grid represents the video layer.
17 Additionally, entries made within the grid represent output pins of the matrix
18 switch. This will become apparent below.

19 The way that the grid is built in this example is that the rendering engine
20 does a traversal operation on the tree 1200. In this particular example, the
21 traversal operation is known as a "depth-first, left-to-right" traversal. This
22 operation will layerize the nodes so that the leftmost track or source has the lowest
23 priority and so on. Doing the above-mentioned traversal on tree 1200 (Fig. 12),
24 the first node encountered is node 1204 which is associated with source A. This is
25 the lowest priority track or source. A first row is defined for the grid and is

1 associated with source A. After the first grid row is defined, a grid entry is made
2 and represents the time period for which source A desires to be routed to the
3 matrix switch's primary output.

4 Fig. 13 shows the state of a grid 1300 after this first processing step.
5 Notice that from time $t = 0-8$, a "0" has been placed in the grid. The "0"
6 represents the output pin of the matrix switch—in this case the primary output.
7 Next, the traversal encounters node 1206 (Fig. 12) which is associated with source
8 B. A second row is thus defined for the grid and is associated with source B.
9 After the second grid row is defined, a grid entry is made and represents the time
10 period for which source B desires to be routed to the matrix switch's primary
11 output.

12 Fig. 14 shows the state of grid 1300 after this second processing step.
13 Notice that from time $t = 4-14$, a "0" has been placed in the grid. Notice at this
14 point that something interesting has occurred which will be resolved below. Each
15 of the layers has a common period of time (i.e. $t = 4-8$) for which it desires to be
16 routed to the matrix switch's primary output. However, because of the nature of
17 the matrix switch, only one input can be routed to the primary output at a time.
18 Next, the traversal encounters node 1208 (Fig. 12) which is associated with source
19 C. In this particular processing example, a rule is defined that sources on tracks
20 are processed before transitions on the tracks are processed because transitions
21 operate on two objects that are beneath them. A third row is thus defined for the
22 grid and is associated with source C. After the third row is defined, a grid entry is
23 made and represents the time period for which source C desires to be routed to the
24 matrix switch's primary output.
25

Fig. 15 shows the state of grid 1300 after this third processing step. Notice that from time $t = 12-16$, a "0" has been placed in the grid. Next, the traversal encounters node 1210 (Fig. 12) which corresponds to the transition. Thus, a fourth row is defined in the grid and is associated with the transition. After the fourth row is defined, a grid entry is made and represents the time period for which the transition desires to be routed to the matrix switch's primary output.

Fig. 16 shows the state of grid 1300 after this fourth processing step. Notice that from time $t = 12-14$, a "0" has been placed in the grid for the transition entry. The transition is a special grid entry. Recall that the transition is programmed to operate on two inputs and provide a single output. Accordingly, starting at the transition entry in the grid and working backward, each of the entries corresponding to the same tree level are examined to ascertain whether they contain entries that indicate that they want to be routed to the output during the same time that the transition is to be routed to the output. If grid entries are found that conflict with the transition's grid entry, the conflicting grid entry is changed to a value to corresponds to an output pin that serves as an input to the transition element 1102 (Fig. 11). This is essentially a redirection operation. In the illustrated grid example, the transition first finds the level that corresponds to source C. This level conflicts with the transition's grid entry for the time period $t = 12-14$. Thus, for this time period, the grid entry for level C is changed to a switch output that corresponds to an input for the transition element. In this example, a "2" is placed in the grid to signify that for this given time period, this input is routed to output pin 2. Similarly, continuing up the grid, the next level that conflicts with the transition's grid entry is the level that corresponds to source B. Thus, for the conflicting time period, the grid entry for level B is changed to a

switch output that corresponds to an input for the transition element. In this example, a "1" is placed in the grid to signify that for this given time period, this input is routed to output pin 1 of the matrix switch.

Fig. 17 shows the state of the grid at this point in the processing. Next, a pruning function is implemented which removes any other lower priority entry that is contending for the output with a higher priority entry. In the example, the portion of A from $t=4-8$ gets removed because the higher priority B wants the output for that time.

Fig. 18 shows the grid with a cross-hatched area that signifies that portion of A's grid entry that has been removed.

At this point, the grid is in a state in which it can be used to program the matrix switch. The left side entries -- A, B, C, and TRANS represent input pin numbers 0, 1, 2, and 3 (as shown) respectively, on the matrix switch shown in Fig. 11. The output pin numbers of the matrix switch are designated at 0, 1, and 2 both on the switch in Fig. 11 and within the grid in Fig. 18. As one proceeds through the grid, starting with source A, the programming of the matrix switch can be ascertained as follows: A is routed to output pin 0 of the matrix switch (the primary output) from $t = 0-4$. From $t = 4-16$, A is not routed to any output pins. From $t = 0-4$, B is not routed to any of the output pins of the matrix switch. From $t = 4-12$, B is routed to the primary output pin 0 of the matrix switch. From $t = 12-14$, B is routed to output pin 1 of the matrix switch. Output pin 1 of the matrix switch corresponds to one of the input pins for the transition element 1102 (Fig. 11). From $t = 14-16$, B is not routed to any of the output pins of the matrix switch. From $t = 0-12$, C is not routed to any of the output pins of the matrix switch. From $t = 12-14$, C is routed to output pin 2 of the matrix switch. Output pin 2 of the

1 matrix switch corresponds to one of the input pins for the transition element 302
2 (Fig. 3). From $t = 12-14$ the transition element (input pin 3) is routed to output pin
3 0. From $t = 14-16$, C is routed to output pin 0 of the matrix switch.

4 As alluded to above, one of the innovative aspects of the matrix switch 308
5 is its ability to seek to any point in a source, without having to process the
6 intervening content serially through the filter. Rather, matrix switch 308 identifies
7 an appropriate transition point and dumps at least a subset of the intervening
8 content, and continues processing from the sought point in the content.

9 The ability of the matrix switch 308 to seek to any point in the media
10 content gives rise to certain performance enhancement heretofore unavailable in
11 computer implemented media processing systems. For example, generation of a
12 filter graph by render engine 222 may take into account certain performance
13 characteristics of the media processing system which will execute the user-defined
14 media processing project. In accordance with this example implementation,
15 render engine 222 may access and analyze the system registry of the operating
16 system, for example, to ascertain the performance characteristics of hardware
17 and/or software elements of the computing system implementing the media
18 processing system, and adjust the filter graph construction to improve the
19 perceived performance of the media processing system by the user. Nonetheless,
20 there will always be a chance that a particular instance of a filter graph will not be
21 able to process the media stream fast enough to provide the desired output at the
22 desired time, i.e., processing of the media stream bogs down leading to delays at
23 the rendering filter. In such a case, matrix switch 308 will recognize that it is not
24 receiving media content at the appropriate project time, and may skip certain
25 sections of the project in an effort to "catch-up" and continue the remainder of the

1 project in real time. According to one implementation, when matrix switch 308
2 detects such a lag in processing, it will analyze the degree of the lag and issue a
3 seek command to the source (through the source processing chain) to a future
4 point in the project, where processing continues without processing any further
5 content prior to the sought point.

6 Thus, for the editing project depicted in Fig. 10, the processing described
7 above first builds a data structure (i.e. data structure 1200 in Fig. 12) that
8 represents the project in hierarchical space, and then uses this data structure to
9 define or create another data structure that can be utilized to program the matrix
10 switch.

11 Fig. 19 is a flow diagram that describes steps in a method in accordance
12 with the described embodiment. The method can be implemented in any suitable
13 hardware, software, firmware, or combination thereof. In the illustrated and
14 described embodiment, the method is implemented in software.

15 Step 1900 provides a matrix switch. An exemplary matrix switch is
16 described above. Step 1902 defines a first data structure that represents the editing
17 project. Any suitable data structure can be used, as will be apparent to those of
18 skill in the art. In the illustrated and described embodiment, the data structure
19 comprises a hierarchical tree structure having nodes that can represent tracks
20 (having one or more sources), composites, transitions and effects. Step 1904
21 processes the first data structure to provide a second data structure that is
22 configured to program the matrix switch. Any suitable data structure can be
23 utilized to implement the second data structure. In the illustrated and described
24 embodiment, a grid structure is utilized. Exemplary processing techniques for
25 processing the first data structure to provide the second data structure are

described above. Step 1906 then uses the second data structure to program the matrix switch.

Example Project with a Transition and an Effect

Consider project 2000 depicted in Fig. 20. In this project there are three tracks, each of which contains a source, i.e. source A, B and C. This project includes an effect applied on source B and a transition between sources B and C. The times are indicated as shown.

As the user creates their project, a data structure representing the project is built. Fig. 21 shows an exemplary data structure in the form of a hierarchical tree 2100 that represents project 2000. There, the data structure includes three tracks, each of which contains one of the sources. The sources are arranged in the tree structure in the order of their priority, starting with the lowest priority source on the left and proceeding to the right. There is an effect (i.e. “Fx”) that is attached to or otherwise associated with source B. Additionally, there is a transition attached to or otherwise associated with source C.

In building the grid for project 2000, the following rule is employed for effects. An effect, in this example, is a one-input/one-output object that is applied to one object—in this case source B. When the effect is inserted into the grid, it looks for any one object beneath it in priority that has a desire to be routed to the primary output of the matrix switch at the same time. When it finds a suitable object, it redirects that object’s output from the matrix switch’s primary output to an output associated with the effect.

As an example, consider Fig. 22 and the grid 2200. At this point in the processing of tree 2100, the rendering engine has incorporated entries in the grid

corresponding to sources A, B and the effect. It has done so by traversing the tree 2100 in the above-described way. In this example, the effect has already looked for an object beneath it in priority that is competing for the primary output of the matrix switch. It found an entry for source B and then redirected B's grid entry to a matrix switch output pin that corresponds to the effect—here output pin 1.

As the render engine 222 completes its traversal of tree 2100, it completes the grid. Fig. 23 shows a completed grid 2200. Processing of the grid after that which is indicated in Fig. 22 takes place substantially as described above with respect to the first example. Summarizing, this processing though: after the effect is entered into the grid and processed as described above, the traversal of tree 2100 next encounters the node associated with source C. Thus, a row is added in the grid for source C and an entry is made to indicate that source C desires the output from $t = 12-16$. Next, the tree traversal encounters the node associated with the transition. Accordingly, a row is added to the grid for the transition and a grid entry is made to indicate that the transition desires the output from $t = 12-14$. Now, as described above, the grid is examined to find two entries, lower in priority than the transition and located at the same tree level as the transition, that compete for the primary output of the matrix switch. Here, those entries correspond to the grid entries for the effect and source C that occur from $t = 12-14$. These grid entries are thus redirected to output pins of the matrix switch 308 that correspond to the transition—here pins 2 and 3 as indicated. Next, the grid is pruned which, in this example, removes a portion of the grid entry corresponding to source A for $t = 4-8$ because of a conflict with the higher-priority entry for source B.

Fig. 24 shows the resultant matrix switch that has been built and configured as the grid was being processed above. At this point, the grid can be used to program the matrix switch. From the grid picture, it is very easy to see how the matrix switch 308 is going to be programmed. Source A will be routed to the matrix switch's primary output (pin 0) from $t = 0-4$; source B will be redirected to output pin 1 (effect) from $t = 4-14$ and the effect on B will be routed to the output pin 0 from $t = 4-12$. From $t = 12-14$, the effect and source C will be routed to output pins corresponding to the transition (pins 2 and 3) and, accordingly, during this time the transition (input pin 4) will be routed to the primary output (output pin 0) of the matrix switch. From $t = 14-16$, source C will be routed to the primary output of the matrix switch.

It will be appreciated that as the software, in this case the render engine 222, traverses the tree structure that represents a project, it also builds the appropriate graphs and adds the appropriate filters and graphs to the matrix switch. Thus, for example, as the render engine 222 encounters a tree node associated with source A, in addition to adding an entry to the appropriate grid, the software builds the appropriate graphs (i.e. collection of linked filters), and associates those filters with an input of the matrix switch. Similarly, when the render engine 222 encounters an effect node in the tree, the software obtains an effect object or filter and associates it with the appropriate output of the matrix switch. Thus, in the above examples, traversal of the tree structure representing the project also enables the software to construct the appropriate graphs and obtain the appropriate objects and associate those items with the appropriate inputs/outputs of the matrix switch 308. Upon completion of the tree traversal and processing of the grid, an

appropriate matrix switch has been constructed, and the programming (i.e. timing) of inputs to outputs for the matrix switch has been completed.

Treatment of “blanks” in a Project

There may be instances in a project when a user leaves a blank in the project time line. During this blank period, no video or audio is scheduled for play.

Fig. 25 shows a project that has such a blank incorporated therein. If there is such a blank left in a project, the software is configured to obtain a “black” source and associate the source with the matrix switch at the appropriate input pin. The grid is then configured when it is built to route the black source to the output at the appropriate times and fade from the black (and silent) source to the next source at the appropriate times. The black source can also be used if there is a transition placed on a source for which there is no additional source from which to transition.

Audio Mixing

In the examples discussed above, sources comprising video streams were discussed. In those examples, at any one time, only two video streams were combined into one video stream. However, each project can, and usually does contain an audio component. Alternately, a project can contain only an audio component. The audio component can typically comprise a number of different audio streams that are combined. The discussion below sets forth but one way of processing and combining audio streams.

In the illustrated example, there is no limit on the number of audio streams that can be combined at any one time.

Suppose, for example, there is an audio project that comprises 5 tracks, A-E. Fig. 26 shows an exemplary project. The shaded portions of each track represent the time during which the track is not playing. So, for example, at $t=0-4$, tracks B, D, and E are mixed together and will play. From $t = 4-10$, tracks A-E are mixed together and will play, and the like.

Fig. 27 shows the grid for this project at 2700. Since we are dealing with this composition now, all of the effects and transitions including the audio mixing are only allowed to affect things in this composition. Thus, there is the concept of a boundary 2702 that prevents any actions or operations in this composition from affecting any other grid entries. Note that there are other entries in the grid and that the presently-illustrated entries represent only those portions of the project that relate to the audio mixing function.

Grid 2700 is essentially set up in a manner similar to that described above with respect to the video projects. That is, for each track, a row is added to the grid and a grid entry is made for the time period during which the source on that track desires to be routed to the primary output of the matrix switch. In the present example, grid entries are made for sources A-E. Next, in the same way that a transition or effect was allocated a row in the grid, a "mix" element is allocated a row in the grid as shown and a grid entry is made to indicate that the mix element desires to be routed to the primary output of the matrix switch for a period of time during which two or more sources compete for the matrix switch's primary output. Note that in this embodiment, allocation of a grid row for the mix element can be implied. Specifically, whereas in the case of a video project,

overlapping sources simply result in playing the higher priority source (unless the user defines a transition between them), in the audio realm, overlapping sources are treated as an implicit request to mix them. Thus, the mix element is allocated a grid row any time there are two or more overlapping sources.

Once the mix element is allocated into the grid, the grid is processed to redirect any conflicting source entries to matrix switch output pins that correspond to the mix element. In the above case, redirection of the grid entries starts with pin 3 and proceeds through to pin 7. The corresponding matrix switch is shown in Fig. 28. Notice that all of the sources are now redirected through the mix element which is a multi-input/one output element. The mix element's output is fed back around and becomes input pin 15 of the matrix switch. All of the programming of the matrix switch is now reflected in the grid 2700. Specifically, for the indicated time period in the grid, each of the sources is routed to the mix element which, in turn, mixes the appropriate audio streams and presents them to the primary output pin 0 of the matrix switch.

Compositions

There are situations that can arise when building an editing project where it would be desirable to apply an effect or a transition on just a subset of a particular project or track. Yet, there is no practicable way to incorporate the desired effect or transition. In the past, attempts to provide added flexibility for editing projects have been made in the form of so called "bounce tracks", as will be appreciated and understood by those of skill in the art. The use of bounce tracks essentially involves processing various video layers (i.e. tracks), writing or moving the processed layers or tracks to another location, and retrieving the processed layers

when later needed for additional processing with other layers or tracks. This type of processing can be slow and inefficient.

To provide added flexibility and efficiency for multi-media editing projects, the notion of a *composite or composition* is introduced. A composite or composition can be considered as a representation of an editing project as a single track. Recall that editing projects can have one or more tracks, and each track can be associated with one or more sources that can have effects applied on them or transitions between them. In addition, compositions can be nested inside one another.

Example Project with Composite

Consider, for example, Fig. 29 which illustrates an exemplary project 2900 having a composition 2902. In this example, composition 2902 comprises sources B and C and a transition between B and C that occurs between $t = 12-14$. This composition is treated as an individual track or layer. Project 2900 also includes a source A, and a transition between source A and composition 2902 at $t = 4-8$. It will be appreciated that compositions can be much more complicated than the illustrated composition, which is provided for exemplary purposes only. Compositions are useful because they allow the grouping of a particular set of operations on one or more tracks. The operation set is performed on the grouping, and does not affect tracks that are not within the grouping. To draw an analogy, a composition is similar in principle to a mathematical parenthesis. Those operations that appear within the parenthesis are carried out in conjunction with those operations that are intended to operate of the subject matter of the

1 parenthesis. The operations within the parenthesis do not affect tracks that do not
2 appear within the parenthesis.

3 In accordance with the processing that is described above in connection
4 with Fig. 19, a first data structure is defined that represents the editing project.
5 Fig. 30 shows an exemplary data structure 3000 in the form of a hierarchical tree
6 structure. In this example, group node 3002 includes two children—track node
7 3004 and composite node 3006. Track node 3004 is associated with source A.
8 Composite node 3006 includes two children—track nodes 3008 and 3010 that are
9 respectively associated with sources B (3008a) and C (3010a). A transition T2
10 (3012) is applied on source C and a transition T1 (3014) is applied on composition
11 3006.

12 Next, data structure 3000 is processed to provide a second data structure
13 that is configured to program the matrix switch. Note that as the data structure is
14 being programmed, a matrix switch is being built and configured at the same time.
15 In this example, the second data structure comprises a grid structure that is
16 assembled in much the same way as was described above. There are, however,
17 some differences and, for purposes of understanding, the complete evolution of the
18 grid structure is described here. In the discussion that follows, the completed
19 matrix switch is shown in Fig. 38.

20 When the rendering engine initiates the depth-first, left-to-right traversal of
21 data structure 3000, the first node it encounters is track node 3004 which is
22 associated with source A. Thus, a first row of the grid is defined and a grid entry
23 is made that represents the time period for which source A desires to be routed to
24 the matrix switch's primary output pin.
25

Fig. 31 shows the state of a grid 3100 after this first processing step. Next the traversal of data structure 3000 encounters the composite node 3006. The composite node is associated with two tracks—track 3008 and track 3010. Track 3008 is associated with source B. Accordingly, a second row of the grid is defined and a grid entry is made that represents the time period for which source B desires to be routed to the matrix switch's primary output pin. Additionally, since B is a member of a composition, meta-information is contained in the grid that indicates that this grid row defines one boundary of the composition. This meta-information is graphically depicted with a bracket that appears to the left of the grid row.

Fig. 32 shows the state of grid 3100 after this processing step. Next, the traversal of data structure 3000 encounters node 3010 which is associated with source C. Thus, a third row of the grid is added and a grid entry is made that represents the time period for which source C desires to be routed to the matrix switch's primary output pin.

Fig. 33 shows the state of grid 3100 after this processing step. Notice that the bracket designating the composition now encompasses the grid row associated with source C. The traversal next encounters node 3012 which is the node associated with the *second* transition T2. Thus, as in the above example, a grid row is added for the transition and a grid entry is made that represents the time period for which the transition desires to be routed to the matrix switch's primary output pin.

Fig. 34 shows the state of grid 3100 after this processing step. Notice that the bracket designating the composition is now completed and encompasses grid row entries that correspond to sources B and C and the transition between them.

Recall from the examples above that a transition, in this example, is programmed to operate on two inputs and provide a single output. In this instance, and because the transition occurs within a composition, the transition is constrained by a rule that does not allow it to operate on any elements outside of the composition. Thus, starting at the transition entry and working backward through the grid, entries at the same tree level and within the composition (as designated by the bracket) are examined to ascertain whether they contain entries that indicate that they want to be routed to the output during the same time that the transition is to be routed to the output. Here, both of the entries for sources B and C have portions that conflict with the transition's entry. Accordingly, those portions of the grid entries for sources B and C are redirected or changed to correspond to output pins that are associated with a transition element that corresponds to transition T2.

Fig. 35 shows the state of grid 3100 after this processing step. The traversal next encounters node 3014 which is the node that is associated with the transition that occurs between source A and composition 2902 (Fig. 29). Processing of this transition is similar to processing of the transition immediately above except for the fact that the transition does not occur within the composition. Because the transition occurs between the composition and another source, one of the inputs for the transition will be the composition, and one of the inputs will be source A (which is outside of the composition). Thus, a grid row is added for this transition and a grid entry is made that represents the time period for which the transition desires to be routed to the matrix switch's primary output pin.

Fig. 36 shows the state of grid 3100 after this processing step. At this point then, the grid is examined for entries that conflict with the entry for transition T1.

One conflicting grid entry is found for the row that corresponds to source B (inside the composition) and one that corresponds to source A (outside the composition). Accordingly, those portions of the grid row that conflict with transition T1 are changed or redirected to have values that are associated with output pins of the matrix switch that are themselves associated with a transition element T1. In this example, redirection causes an entry of “3” and “4” to be inserted as shown.

Fig. 37 shows the state of grid 3100 after this processing step. If necessary, a pruning operation would further ensure that the grid has no competing entries for the primary output of the matrix switch. The associated input pin numbers of the matrix switch are shown to the left of grid 3100.

Fig. 38 shows a suitably configured matrix switch that has been build in accordance with the processing described above. Recall that, as data structure 3000 (Fig. 30) is processed by the rendering engine, a matrix switch is built and configured in parallel with the building and processing of the grid structure that is utilized to program the matrix switch. From the matrix switch and grid 3100 of Fig. 37, the programming of the switch can be easily ascertained.

Fig. 38a shows an exemplary data structure that represents a project that illustrates the usefulness of composites. In this example, the project can mathematically be represented as follows:

$$(Fx\text{-noisy } (A \text{ Tx-Blend } B)) \text{ Tx-Blend } C$$

Here, an effect (noisy) is applied to A blended with B, the result of which is applied to a blend with C. The composite in this example allows the grouping of the things beneath it so that the effect (noisy), when it is applied, is applied to

1 everything that is beneath it. Notice that without the composite node, there is no
2 node where an effect can be applied that will affect (A Tx-Blend B). Hence, in
3 this example, operations that appear within the parenthesis are carried out on
4 tracks that appear within the parenthesis. Those operations do not affect tracks
5 that are not within the parenthesis.

6 Fig. 39 is a flow diagram that described steps in a method in accordance
7 with one embodiment. The method can be implemented in any suitable hardware,
8 software, firmware, or combination thereof. In the presently-described example,
9 the method is implemented in software.

10 Step 3900 defines a multimedia editing project that includes at least one
11 composite. The composite represents multiple tracks as a single track for purposes
12 of the processing described just below. It is important to note that, in the
13 processing described just below, and because of the use of composites, the extra
14 processing that is required by bounce tracks is avoided (i.e. operating on two
15 tracks, moving the operation result to another location, and retrieving the
16 operation result when later needed). This reduces the processing time that is
17 required to render a multi-media project. Step 3902 defines a first data structure
18 that represents the editing project. Any suitable data structure can be utilized. In
19 the present example, a data structure in the form of a hierarchical tree is utilized.
20 An exemplary tree is shown in Fig. 30. Step 3904 processes the first data structure
21 to provide a second data structure that is configured to program a matrix switch.
22 In the illustrated example, the second data structure comprises a grid structure.
23 Exemplary processing is described in the context of Figs. 30-37. Step 3906 then
24 programs the matrix switch using the second data structure.
25

Source Combining

Having introduced the various architectural and implementation elements of the present invention, above, attention is now drawn to Figs. 40-43, wherein another aspect of the illustrated embodiment is presented. As introduced above, the opening and processing of media represents consumption of memory and processing resources. Thus, performance improvements may be achieved by reducing the number of times a source is accessed. Thus, a method is presented in accordance with one aspect of the present invention that serves to reduce the number of times a source is accessed, e.g., a method of source combining. It is to be appreciated, however, that the following is but one example implementation of the broader inventive concept of reducing the number of times a source need be accessed during execution of a development project. Alternative methods of source combining of greater or lesser complexity may well be used within the spirit and scope of the present invention. Indeed, such alternative methods are anticipated within the scope of the present invention.

Fig. 40 illustrates an example method of generating a filter graph, in accordance with one aspect of the present invention. As shown, method 4000 begins with block 4002, wherein render engine 222 receives an indication to generate a development project. According to one implementation, as discussed above, render engine 222 receives the indication from a higher-level application 216, e.g., media processing application 224, to assist a user in generating a processing project (e.g., a media processing project).

In block 4004, render engine 222 identifies the number and nature of the media sources within the user-defined processing project, in preparation for generating a filter graph representation of the processing project. As introduced

1 above, for each of the identified sources, render engine 222 determines the
2 necessary transform filters 306 required to pre-process the source (i.e., source
3 chain), preparing the processing chain for presentation to the matrix switch filter
4 308 and one or more transition/effect filters 306. Unlike conventional
5 implementations, which would proceed to generate the entire filter graph in
6 preparation for execution of the processing project, render engine 222 generates a
7 list of sources and when they are required in the filter graph. According to one
8 implementation, the list is referred to as a reuse list, and is maintained within
9 render engine 222. An example of a data structure comprising a reuse list is
10 presented with reference to Fig. 41.

11 Turning briefly to Fig. 41, a graphical illustration of an example data
12 structure comprising a source reuse list is presented. As shown, the reuse list 4100
13 is comprised of a number of information fields, e.g., 4102-4110 which detail, in
14 part, the relationship between clips in a track. More particularly, the reuse list
15 4100 is shown comprising a track identification field 4102, a source identification
16 field 4104, a project time field 4106 and a source time field 4108.

17 Upon identifying a project source and the associated filters required for pre-
18 processing the source (i.e., the source chain), render engine 222 assigns each track
19 an identifier which uniquely identifies the source track within the context of the
20 filter graph. In this regard, reuse list 4100 includes a field 4102 which maintains a
21 list of tracks utilized in the associated project. In accordance with the illustrated
22 example paradigm of the media processing system, the track identifier is utilized
23 to represent a media clip from a given source.

24 The source identifier field 4104 contains information denoting the project
25 source associated with a particular track identifier. In this regard, the source

1 identifier field 4104 may well contain a file name, a file handle, or any other
2 suitable source identifier.

3 The project time field 4106 denotes at what point during project execution
4 the media clip is required. The source time field 4108 denotes what portion of the
5 source file is required to support execution of the processing project. It should be
6 appreciated that a user may well utilize the whole source file or any part thereof,
7 as defined by the processing project.

8 In accordance with the illustrated example implementation of Fig. 41, two
9 tracks are depicted 4110 and 4112. As shown, each of the tracks represent media
10 from a common source (e.g., source ID 4213) and, the source media clips are
11 adjacent to one another in the project (e.g., project time 4106) as well as within the
12 source file (e.g., source time 4108). As will be developed more fully below,
13 source clips may well be combined in certain situations into a single clip, as is
14 represented by track 4114 in Fig. 41. It is to be appreciated that, although depicted
15 as a two-dimensional data structure, reuse lists of greater or lesser complexity may
16 well be substituted without deviating from the spirit and scope of the present
17 invention.

18 Returning to Fig 40 and, in particular, block 4006, render engine 222
19 reduces the number of source accesses where possible, in accordance with one
20 aspect of the present invention. More particularly, render engine 222 analyzes the
21 reuse list 4100 to identify opportunities to reduce the number of source accesses
22 by combining source clips which meet certain criteria. According to one
23 implementation, the criteria used by render engine 222 include one or more of (1)
24 the source clips must occur next to one another in the project, (2) the clips appear
25 next to one another in the source, and (3) the clips must share a common pre-

processing source chain (i.e., must require the same pre-processing (e.g., same processing rate, etc.)). If this criteria is met, render engine 222 may combine the clips into a single clip. More specifically, render engine 222 modifies the reuse list 4100 (Fig. 41) to replace the multiple source accesses (4110, 4112) with a single source access 4114 representing both source accesses as a single access. It is to be appreciated that removing a source access improves filter graph performance and, accordingly, the perceived performance of the development system by the user.

In block 4008, once render engine 222 has reduced the number of source file accesses (block 4006), render engine 222 dynamically generates and manages the filter graph to support execution of the development project. In accordance with one aspect of the present invention, render engine 222 invokes only those source chains associated with sources that are necessary to support the current and/or impending execution of the filter graph. It is to be appreciated that by not opening each of the chains of a processing project, render engine 222 reduces the amount of memory required to build the filter graph, thereby reducing the amount of memory required to complete execution of the project.

Turning to **Fig. 42**, an example method for source combining is presented, in accordance with one aspect of the present invention. As shown, the method begins with block 4202, wherein render engine 222 identifies adjacent clips from a common source, i.e., project aligned clips. As introduced above, render engine 222 analyzes the reuse list 4100 to identify all of the clips associated with a particular track. For each of the adjacent source clips within a track, render engine 222 determines whether the clips are adjacent to one another with respect to their source time, block 4204. That is, identifying that the source clips are adjacent to

09734529-120600

one another by project time (e.g., occurring within the same track), render engine 222 next determines whether the clips are adjacent to one another in the source file, i.e., whether the clips are source aligned. According to one implementation, if the source clips are project and source time aligned, render engine 222 determines whether the source clips share a common preprocessing source chain (i.e., the require the same pre-processing).

If the source clips are source aligned, render engine 222 next determines whether the clips require unique pre-processing (e.g., decoding, frame rate conversion, sizing, etc.), block 4206. If unique pre-processing is required (block 4206), or the adjacent project clips are not source aligned (block 4204), the source clips will require an independent source chain of filters and, thus cannot be source combined. Accordingly, the source clips are accessed and processed independently, block 4208.

If, in block 4206, render engine 222 determines that the clips are source aligned and each share common pre-processing requirements, render engine 222 combines the adjacent source clips into a single clip and updates the reuse list 4100 accordingly, block 4210 (e.g., clip 4114 of Fig. 41). The combined source clip is representative of each of the otherwise individual clips, while requiring a single source processing chain and, thus, a single source access. In block 4212, render engine 222 determines whether all of the clips of the reuse list has been analyzed and, if not, the process continues in an iterative manner with block 4202 until the entire reuse list has been analyzed and appropriate source clips combined.

Fig. 43 graphically illustrates the source combining aspect of the present invention. As shown, Fig. 43 illustrates a project 4300 of two tracks of clips (e.g., clips 4302-4318, and 4320-4324) separated by a transition 4318. In accordance

1 with the source combining aspect of the present invention, introduced above,
2 render engine 222 analyzes a reuse list 4100 representation of project 4300 to
3 identify source clips which may be combined. As introduced above, render engine
4 222 will combine source clips which are project and source time aligned, and
5 which do not require unique pre-processing.

6 In the illustrated example of Fig. 43, clips 4302 and 4304 are project
7 aligned and source aligned. If such clips do not require independent pre-
8 processing, they are combined into a single clip 4326 by render engine 222. Note
9 that although clips 4304 and 4308 are project aligned, they are not source aligned
10 (e.g., the media end time (9) of clip 4304 does not abut the media start time (10) of
11 clip 4308), i.e., there is a gap of one elemental unit (e.g., a second of time). The
12 process of source combining is performed for other clips in the development
13 project, reducing the total number of source clips from eleven to six in the
14 illustrated example of Fig. 43. Thus, it is to be appreciated that the source
15 combining aspect of the present invention represents another feature which
16 reduces the computational and memory requirements necessary to support even
17 the most complex development projects.

18 19 **Source Filter Reuse**

20 As introduced above, conventional media processing systems may generate
21 an individual thread each time content was required from a source, even if the
22 source had been accessed earlier. This redundant loading/unloading of a source is
23 computationally expensive, and consumes precious memory resources. Extending
24 the concept of source combining introduced above, a filter and related methods for
25

1 sharing a common source and source filter among multiple processing threads will
2 now be introduced, in accordance with yet another aspect of the present invention.

3 Turning to **Fig. 44**, a block diagram of an example filter graph 4400 is
4 presented incorporating a segment filter 4406 which, as will be shown,
5 dynamically couples a source filter to one or more processing chain in accordance
6 with the teachings of the present invention. In accordance with the illustrated
7 example of Fig. 44, filter graph 4400 is depicted comprising a source 4402, one or
8 more pre-processing transform filters 4404, a segment filter 4406 and one or more
9 pre-processing transform filter(s) 4408A-N, each coupled to a matrix switch 308
10 and rendering filter(s) 4410, 4412, respectively.

11 As used herein, segment filter 4406 is designed to sit between a source
12 filter and matrix switch 308 to provide multiple processing chains with source
13 content from a single source, where it is impossible to combine the source clips (as
14 introduced above). Render engine 222 invokes an instance of the segment filter
15 4406 after the greatest common pre-processing filter 4404 for each of the chains.
16 That is, each of the processing chains may require the source content in a different
17 format (e.g., size, frame.rate, decode format, etc.). To the extent that the chains
18 share common pre-processing attributes, those filter(s) (4404) are placed before
19 the segment filter 4406 where practicable. In many instances, none of the chains
20 share common pre-processing and the pre-processing filter(s) merely comprise the
21 source filter.

22 The segment filter 4406 acts as a controller, or throttle for the source,
23 instructing the source filter to deliver content from source 4402 at select times.
24 According to one implementation, the segment filter 4406 is, in turn, controlled by
25 the render engine 222 and/or the matrix switch filter 308 to provide select content

1 at select times on select inputs of the matrix switch filter 308. According to one
2 implementation, the segment filter 4406 issues a "seek" command to the source
3 filter to request particular content from the source. The source filter then delivers
4 the requested content through the segment filter 4406 and appropriate pre-
5 processing filter(s) 4408A-N to deliver the desired content to the requesting matrix
6 switch 308 to support processing of the development project.

7 As introduced above, render engine 222 is responsive to higher-level user
8 interfaces, e.g., applications 224. In this regard, it is possible that the filter graph
9 will receive user-commands while the filter graph is executing the development
10 project. In accordance with the media processing system paradigm, for example,
11 it is foreseeable that a user-invoked seek will be received by the filter graph during
12 execution of the development project. Such user defined commands are typically
13 serialized with commands issued by filters within the filter graph during the
14 normal course of execution. In accordance with the illustrated example
15 implementation, where matrix switch 308 "throttles" execution of the filter graph,
16 matrix switch 308 issues a seek command of its own to the source filter, requesting
17 the information desired by the user. According to an alternate embodiment, seeks
18 received from a higher-level application (and, therefore, representative of a user
19 command) are afforded a higher priority within the filter graph. In such an
20 implementation, all segment filters 4406 residing within the filter graph are also
21 notified of such high-priority seeks, so that they can identify what content they
22 will be required to provide next and, therefore, issue a revised seek command of
23 their own.

24 The remaining pre-processing transform filter(s) 4408A-N, matrix switch
25 filter(s) 308 and rendering filter(s) 4410 each function as described above.

Turning now to Fig. 45, an example method for generating a filter graph is presented incorporating the teachings of the present invention. More particularly, the method of Fig. 45 is similar to the method of Fig. 42 wherein render engine 222 attempted source combining of source clips, which were not project and source time aligned, or which required unique pre-processing of some sort. In Fig. 42, however, if the source clips were not source time aligned (4204) and/or the clips required separate pre-processing (block 4206), each clip was assigned to a separate processing chain. In Fig. 45, however, this problem is resolved with introduction of a segment filter 4406.

More specifically, with reference to Fig. 45, render engine 222 identifies multiple source clips from a common source which are not source time aligned, block 4204 and/or require separate pre-processing filter(s), block 4206. Render engine 222 generates a segment filter 4406 for the filter graph to reuse the source and at least the source filter, block 4502. That is, the render engine 222 inserts a segment filter 4406 between the source filter and one or more processing chains to selectively provide otherwise disparate source clips from a single source. But for use of the segment filter 4406 in the filter graph, the method 4500 of Fig. 45 executes in a fashion similar to Fig. 42, above.

Turning now to Fig. 46, a flow chart of an example method of segment filter operation is presented, in accordance with one example implementation of this aspect of the present invention. In accordance with the illustrated example embodiment of Fig. 46, the method begins in block 4602, wherein the segment filter 4406 seeks the source to the place that source data is first needed. As introduced above, segment filter 4406 receives a request for source content from matrix switch filter 308. It should be appreciated that insofar as segment filter

1 4406 may well support a plurality of processing chains coupled to a plurality of
2 matrix switch filters 308, a number of such requests may be received
3 simultaneously. According to one implementation, each of the matrix switch
4 filters 308 assigns a priority to the request for source content, wherein the priority
5 of the request changes as the time the content is needed draws near. According to
6 an alternate implementation, render engine 222 determines *a priori* whether source
7 content will be required simultaneously and, if so, provides a separate source chain
8 to accommodate such simultaneous content requests, thereby eliminating the
9 situation of the segment filter 4406 receiving simultaneous requests.

10 In block 4604, the source filter retrieves the requested content and passes
11 the data to the switch until some sort of indication is received that the end of
12 content has been received (e.g., an end-of-stream (EOS) indication, an application
13 interrupt, etc.). As introduced above, an application interrupt may be issued when
14 a user, through a user interface (e.g., media control application 224), wants to seek
15 to a certain point in the development project.

16 In block 4606, segment filter 4406 determines whether an EOS or an
17 application interrupt is received. If not, the process continues with block 4604. If
18 so, segment filter 4406 identifies the next required segment and when it will be
19 required, given the current seek location received from the matrix switch filter
20 308. Based, at least in part, on the current seek location, segment filter 4406
21 determines whether more segments of the source are required, block 4610. As
22 introduced above, if a user-defined seek command is issued, it may be issued to a
23 location in the development project where no further content is required from a
24 particular source. Thus, segment filter 4406 determines whether additional
25 segments are required in block 4610.

1 If no further segments are required along one of the processing chains
2 leading from the segment filter, render engine 222 may remove (at least
3 temporarily) that chain from the filter graph to free memory space and a matrix
4 switch filter input for other processing chains, block 4612.

5 If, in block 4610 further segments are required, segment filter 4406 issues a
6 seek instruction directing the source filter to retrieve and deliver the next segment,
7 in accordance with the matrix switch filter instructions, block 4614. This process
8 continues in an iterative fashion with block 4604.

9 10 **Shared Parser**

11 Having introduced a couple of techniques for reducing the need to open
12 multiple instances of a source file to satisfy content delivery within a single
13 subsystem (e.g., the video subsystem), a system and method will now be
14 introduced which reduces the need to open multiple instances of a source file to
15 satisfy content delivery between separate media processing subsystems. More
16 particularly, a media parser filter is introduced which enables a video processing
17 subsystem of the filter graph (e.g., a video processing chain) and an audio
18 processing subsystem of the filter graph (e.g., an audio processing chain) to share
19 a common instance of a source, in accordance with one aspect of the present
20 invention. As introduced above, a parser is a media processing software object
21 that parses received media content into its constituent elements. According to one
22 implementation, for example, the parser object received multimedia content and
23 separates the multimedia content into its constituent elements of audio content and
24 video content for delivery to appropriate media processing chains.

09734529-120600

Fig. 47 illustrates a block diagram of an example media parser object, according to one embodiment of the present invention. In accordance with the illustrated example embodiment of Fig. 47, a block diagram of an example media parser object 4700 is presented, according to one embodiment. As shown, the media parser object 4700 is depicted comprising one input 4702 and one or more outputs 4704 and 4706, respectively. Input 4702 is dynamically coupled by render engine 222 to a source filter, while outputs 4704 and 4706 are each coupled to a media processing subsystem (i.e., processing filter chain). According to one example implementation, the number of outputs required of a parser is determined by the number of media processing subsystems that require content from a common instance of a source. In accordance with the illustrated example implementation, parser 4700 is configured (i.e., by render engine) with two outputs, i.e., one to a video processing subsystem and one to an audio processing subsystem. It is to be appreciated, however, that parser objects may well be used with only one output, i.e., to separate media content into its constituent elements, provide one of the elements to a requesting media processing subsystem, while discarding the other element(s).

As introduced above, the media processing subsystems are a chain of software objects (e.g., filters) assembled by the render engine to process media content. Thus, in this example, the parser filter is coupled to a video source filter chain with one output (e.g., 4704), and to an audio source filter chain with another output (e.g., 4706). It is important to note, as introduced above, that the media parser filter 4700, like other objects in the filter graph, may well be implemented as instances of a software object, e.g., a component object model (COM) object.

1 According to the teachings of the present invention, media parser object
2 4700 receives media content from a source filter via input 4702, splits the media
3 content into its constituent elements, and passes the separate media content to the
4 appropriate media processing subsystems via outputs 4704 and 4706. Any number
5 of techniques may be employed by media parser filter 4700 to identify and parse
6 the media content types. According to one implementation, for example, media
7 parser filter 4700 receives the media content in accordance with one of a number
8 of well known media communication standards, which reserves particular space
9 within the communication packets for each of the different media content types.
10 In such an implementation, media parser filter 4700 separates the different data
11 portions of the received packets to provide the appropriate media content to the
12 appropriate outputs 4704, 4706. It is to be appreciated that alternate techniques
13 may well be employed within the scope and spirit of the present invention.

14 Parser object 4700 retrieves media content from a source via the source
15 filter in response to requests for content received from the one or more media
16 processing subsystems. According to one implementation, parser object 4700
17 serializes requests received from the media processing subsystems and satisfies
18 such requests in the order in which they were received. According to another
19 implementation, parser object 4700 only responds to requests for content received
20 from one of a potential plurality of media processing subsystems, i.e., from one of
21 its outputs. In such a system, requests received from a different media processing
22 subsystem are ignored.

23 **Fig. 48** illustrates a flow chart of an example method for building a filter
24 graph, according to one aspect of the present invention. In accordance with the
25 illustrated example embodiment of Fig. 48, the method begins with block 4802

1 wherein render engine 222 identifies multiple requests or accesses to a common
2 source. In block 4802, render engine 222 determines whether the multiple
3 requests are from a single media processing subsystem, or from multiple media
4 processing subsystems. If the multiple requests stem from within a single media
5 processing subsystem, render engine 222 invokes an instance of the segment
6 object 4406, described above with reference to Fig. 45, to satisfy these multiple
7 requests, block 4806. It is important to note that render engine 222 may also
8 employ an instance of a parser object 4700 in the source processing chain to parse
9 the desired media content type (i.e., audio, video, etc.) for use by the downstream
10 processing chain. That is, unless the source contains media of only a single type
11 (e.g., a video-only media file), render engine will invoke an instance of a parser
12 object 4700 in conjunction with a segment object 4406 in a processing chain.

13 If, in block 4804 render engine 222 determines that the multiple requests
14 stem from more than one media processing subsystem, render engine 222 invokes
15 an instance of media parser filter 4700 with multiple outputs to couple each of the
16 media processing subsystems to a single instance of the source filter, block 4808.
17 More specifically, render engine 222 invokes an instance of the media parser filter
18 4700, assigning one of the outputs 4704 to the video processing subsystem, the
19 other outputs 4706 to the audio processing subsystem, while input 4702 is coupled
20 to the source filter.

21 In block 4810, media parser filter 4700 selectively provides media content
22 to the media processing subsystems in response to received requests for such
23 content until an indication of completion is received. As introduced above, the
24 requests for such media content may well be received from matrix switch filter
25 308, a higher-level application (e.g., 224) via render engine 222, and the like.

09734569 120600

1 In block 4812, upon receiving an indication of completion from one of the
2 media processing subsystems, media parser object 4700 has to make a
3 determination of whether to unload the source filter. The obvious problem arises
4 if the media parser object 4700 were to unload the source filter before all of the
5 media processing subsystems had completed processing media from the source. In
6 accordance with one implementation, render engine 222 identifies a priority
7 subsystem (e.g., the video processing subsystem) to control the dissolution of the
8 source filter. Thus, the decision in block 4812 is to authenticate whether the
9 indication of completion was received from the priority subsystem. If so, the
10 source filter may well be unloaded (at least temporarily), block 4814. If not,
11 processing continues with block 4810. In an alternate implementation, upon
12 receiving an indication of completion from one of the processing chains, parser
13 object 4700 queries the other processing chain to determine whether it requires
14 further content from the source. Such requests may well be directed to a matrix
15 switch filter 308, if coupled to the media processing chain. If it is determined that
16 one or more of the media processing chains still requires content from the source,
17 the source chain is left untouched. If all media processing subsystems have
18 exhausted their need for content from the source, parser object 4700 informs
19 render engine 222, which may remove the source processing chain from the active
20 filter graph. According to one implementation, if render engine 222 determines
21 that the source processing chain may be used subsequently, i.e., in the current or
22 future filter graphs, it may cache the source processing chain whole, for rapid
23 integration in a subsequent filter graph.

24 **Fig. 49** graphically illustrates an example filter graph sharing a parser
25 object 4700 between different media processing sub-systems, according to one

embodiment of the present invention. In accordance with the illustrated example embodiment of Fig. 49, a media parser object 4700 is depicted providing each of two separate media processing subsystems with media content from a common instance of a source filter 4902. More specifically, filter graph 4900 is presented comprising a single source filter 4902 and two media processing subsystems 4904 and 4908 each coupled to the single source filter 4902 via media parser filter 4700. Each of the media processing subsystems are recursively coupled to one or more effects and/or transition filters (not shown) through matrix switch filter 308, which passes the final composite to their respective render filters 4906 and 4910, respectively. It is to be appreciated, however, that media parser filter 4700 may well be beneficially utilized in filter graphs that do not contain the innovative matrix switch filter 308, or to provide content to a single media processing subsystem without deviating from the spirit and/or scope of the present invention. Accordingly, Fig. 49 also depicts a source filter chain with a source 4912 coupled to a rendering filter through a media processing chain incorporating a parser object 4700, and without the need of the innovative matrix switch filter 308.

According to one embodiment, the video processing subsystem (e.g., 4904) is selected as the priority subsystem. In accordance with such an embodiment completion indications received from the audio processing subsystem (e.g., 4908) alone, go unanswered by media parser object 4700. It is not until the video processing subsystem (e.g., 4904) issues a completion indication, having verified that the audio processing subsystem is also finished accessing the source, that media parser filter 4700 considers removing the source filter 4902 and the associated media processing subsystems 4904, 4908 from the active filter graph. Alternatively, as introduced above, parser object 4700 verifies that all media

1 processing subsystems have completed processing from a source filter before it
2 may be considered for removal from the active filter graph. Thus, it is to be
3 appreciated that use of the innovative media parser filter 4700 provides media
4 content to each of two or more media processing subsystems from a single
5 instance of a source filter, thereby reducing the number of instances of open source
6 files and providing associated performance improvements.

7 Although the invention has been described in language specific to structural
8 features and/or methodological steps, it is to be understood that the invention
9 defined in the appended claims is not necessarily limited to the specific features or
10 steps described. Rather, the specific features and steps are disclosed as preferred
11 forms of implementing the claimed invention.